Last modified: September 19, 2024. Link to latest version.

Kimberly J. Wilber

kimmy@kjwilber.org
http://kjwilber.org

Research engineer with an academic computer vision background. She / Her

Technical Skills

Computer vision and machine learning. I use applied research to empower users and the societies we live in. Fluent in scientific Python tools (TensorFlow, PyTorch, scikit-learn, matplotlib, numpy/scipy). Uses OpenCV, dlib, OpenGL, and Apache Flume in C environments.

Strong general proficiency with POSIX and CLI tools on Linux and Mac OS. Writes side projects in Torch7/Lua, NodeJS, and nim.
Sysadmin. Over 20 years of experience administering Debian and Ubuntu server clusters across various hosting environments (AWS EC2, GCP, and on-prem) using tools like Ansible and Docker.

Full-stack. Frequently works with projects spanning low-level UNIX/C to high level Python to frontend (HTML/Javascript, Coffeescript, Livescript, in Vue) in both SQL database and MapReduce environments.
Open-source. Contributor to projects including Racket and node.js – ask me about the `clearTimeout` nodejs bug I found and fixed in 2010!

Education

2014–2018 **Ph.D. in Computer Science, Cornell Tech**
Supported by the National Science Foundation Graduate Research Fellowship (NSF GRFP)

2013–2014 Graduate studies at University of California, San Diego
Transferred to Cornell to follow my advisor, Dr. Serge Belongie

2009–2013 **Bachelor of Innovation in Computer Science, University of Colorado Colorado Springs**
Supported by the Kane Family Foundation Scholarship

Professional Experience

2013–Present **Volunteer publication assistant, Computer Vision Foundation**
- I help the CV Foundation process papers for the CVF Open Access Archive, a complete repository of CVPR, ICCV, and WACV papers.
- To help prepare conference proceedings, I built a pipeline that processes over 2,700 camera-ready papers provided by authors for each conference.
- Every machine learning researcher who reads open-access CVPR papers in the last 11 years has touched PDF files generated by the infrastructure I maintain.
- My pipeline adds page numbers, an attribution banner, and proper PDF metadata, with automatic checks for common quality problems (figures overlapping margins, etc).

2018–2024 **Software Engineer, Google AI, New York, NY**
- Implemented real-time, on-device monocular depth estimation models for Project Guideline, a tool that helps visually impaired runners exercise independently. My models improved Google's outdoor depth estimation capabilities.
- Helped create SANPO, a 3D panoptic video dataset for the computer vision community. My role was to create large-scale data processing infrastructure and to shepherd the collection of 3D depth data.
- Created and refactored internal debug tools for *RankEmbed*, one of Google Web Search's first forays into NLP embeddings for web search. My tools helped engineers demo the fledgling product to executives to build support.
- Built a reputation as a strong visual communicator – others relied on me to produce publication-ready figures and graphics for papers and presentations authored during my tenure.

2014–2018 **Research Assistant, Cornell University, Cornell Tech NYC**
- Conducted computer vision research on perceptual similarity, crowdsourcing, and object recognition.
- Helped establish and maintain the new vision group's presence at Cornell.
- Served as TA for classes including four semesters of "CS5785 Modern Analytics."

2017 **Summer Intern, Google Photos Team, Mountain View, CA**
- Implemented and tested tools to make it easier for ML engineers to prototype UI interactions.
- These tools helped other engineers decide what features to include in Google Photos Sharing Suggestions.

2016 **Summer Intern, Adobe Research, San Jose, CA**
- Curated BAM, one of the first large-scale collections of professional commercial artwork, intended for ML object classification and emotion understanding.
- Built a data loader system in Python and Redis to quickly analyze millions of images for ML training and inference workloads, speeding up training by 5×.

2014 **Summer Intern, Dropbox Photos Team, San Francisco, CA**
- Conducted product-focused computer vision research.
- Introduced our team to more efficient tools and technologies.
- Maintained a computer vision evaluation and experimentation pipeline.

2013 **Research Assistant, University of California, San Diego, CA**

2012–2013 **Software Engineer, Securics, Inc., Colorado Springs, CO**

2009–2013 **Assistant Researcher, Vision and Security Technology (VAST) Laboratory at UCCS, CO**

2011 **Summer Researcher, NSF REU Program, University of Colorado Colorado Springs, CO**

2009–2010 **NSF RAHSS High School Intern, Securics, Inc., Colorado Springs, CO**

Publications

Note that some work before 2018 is published under a previous name. All titles are clickable.

2024 *PolyMaX: General Dense Prediction with Mask Transformer*
Xuan Yang; Liangzhe Yuan; **Kimberly Wilber**; Astuti Sharma; Xiuye Gu; Siyuan Qiao; Stephanie Debats; Huisheng Wang; Hartwig Adam; Mikhail Sirotenko; Liang-Chieh Chen. *Winter Conference on Applications of Computer Vision (WACV 2024)*

2023 *SANPO: A Scene Understanding, Accessibility, Navigation, Pathfinding, Obstacle Avoidance Dataset*
Sagar M. Waghmare; **Kimberly Wilber**; Dave Hawkey; Xuan Yang; Matthew Wilson; Stephanie Debats; Cattalyya Nuengsigkapian; Astuti Sharma; Lars Pandikow; Huisheng Wang; Hartwig Adam; Mikhail Sirotenko. *ArXiv*

2022 *On Label Granularity and Object Localization*
Elijah Cole; **Kimberly Wilber**; Grant Van Horn; Xuan S. Yang; Marco Fornoni; Pietro Perona; Serge Belongie; Andrew G. Howard; Oisin Mac Aodha. *European Conference on Computer Vision (ECCV 2022)*

2022 *Exploring Fine-Grained Audiovisual Categorization with the SSW60 Dataset*
Grant Van Horn; Rui Qian; **Kimberly Wilber**; Hartwig Adam; Oisin Mac Aodha; Serge Belongie. *European Conference on Computer Vision (ECCV 2022)*

2022 *When Does Contrastive Visual Representation Learning Work?*
Elijah Cole; Xuan Yang; **Kimberly Wilber**; Oisin Mac Aodha; Serge Belongie. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2022)*

2021 *Benchmarking Representation Learning for Natural World Image Collections*
Grant Van Horn; Elijah Cole; Sara Beery; **Kimberly Wilber**; Serge Belongie; Oisin Mac Aodha. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2021)*

2021 *On the Reproducibility of Neural Network Predictions*
Srinadh Bhojanapalli; **Kimberly Wilber**; Andreas Veit; Ankit Rawat; Seungyeon Kim; Aditya Menon; Sanjiv Kumar. *ArXiv*

2021 *Bridging the Gap Between Object Detection and User Intent via Query-Modulation*
Marco Fornoni; Chaochao Yan; Liangchen Luo; **Kimberly Wilber**; Alex Stark; Yin Cui; Boqing Gong; Andrew Howard. *ArXiv*

2020 *Improving Calibration in Deep Metric Learning With Cross-Example Softmax*
Andreas Veit; **Kimberly Wilber**. *ArXiv*

2019 *Understanding Image Quality and Trust in Peer-to-Peer Marketplaces*
Xiao Ma; Lina Mezghani; **Kimberly Wilber**; Hui Hong; Robinson Piramuthu; Mor Naaman; Serge Belongie. *Winter Conference on Applications of Computer Vision (WACV 2019)*

2018 *Learning perceptual similarity from crowds and machines*
M. Wilber. *PhD Thesis, Cornell University, Ithaca, NY. Advised by Serge Belongie.*

2018 *Learning from Multi-domain Artistic Images for Arbitrary Style Transfer*
Zheng Xu; **M. Wilber**; Chen Fang; Aaron Hertzmann; Hailin Jin. *ACM/Eurographics Expressive Symposium*

2017 *BAM! The Behance Artistic Media Dataset for Recognition Beyond Photography*
M. Wilber; Chen Fang; Hailin Jin; Aaron Hertzmann; John Collomosse; Serge Belongie. *International Conference on Computer Vision (ICCV 2017)*

2017 *Sketching with Style: Visual Search with Sketches and Aesthetic Context*
John Collomosse; Tu Bui; **M. Wilber**; Chen Fang; Hailin Jin. *International Conference on Computer Vision (ICCV 2017)*

2017 *Crowd Research: Open and Scalable University Laboratories*
Rajan Vaish; Snehalkumar (Neil) S. Gaikwad; Geza Kovacs; Andreas Veit; Ranjay Krishna; Imanol Arrieta Ibarra; Camelia Simoiu; **M. Wilber**; Serge Belongie; Sharad Goel; James Davis; Michael S. Bernstein. *User Interface Software and Technology Symposium (UIST 2017)*

2016 *Residual Networks Behave Like Ensembles of Relatively Shallow Networks*
Andreas Veit; **M. Wilber**; Serge Belongie. *Neural information processing systems (NIPS 2016)*

2016 *Training and investigating Residual Nets*
Sam Gross; **M. Wilber**. *Tech report (Torch blog)*

2016 *Can we still avoid automatic face detection?*
M. Wilber; Vitaly Shmatikov; Serge Belongie. *Winter Conference on Applications of Computer Vision (WACV 2016)*

2015 *Learning Concept Embeddings with Combined Human-Machine Expertise*
M. Wilber; Iljung Sam Kwak; Serge Belongie. *International Conference on Computer Vision (ICCV 2015)*

2015 *On Optimizing Human-Machine Task Assignments*
Andreas Veit; **M. Wilber**; Rajan Vaish; Serge Belongie; James Davis; et. al.. *AAAI Conference on Human Computation and Crowdsourcing Work-in-Progress session (HCOMP 2015 WIP)*

2015 *Image Representations and New Domains in Neural Image Captioning*
Jack Hessel; Nicolas Savva; **M. Wilber**. *Workshop on Vision and Language Integration (VL 2015)*

2014 *Cost-Effective HITs for Relative Similarity Comparisons*
M. Wilber; Iljung Sam Kwak; Serge Belongie. *AAAI Conference on Human Computation and Crowdsourcing (HCOMP 2014)*

2014 *Exemplar Codes: An Accurate and Efficient Mid-Level Representation for Big Vision Problems*
Ethan Rudd; **M. Wilber**; Terry Boult. *Computer Vision and Pattern Recognition BigVision workshop (CVPR 2014)*

2014 *Exemplar Codes for Facial Attributes and Tattoo Recognition*
M. Wilber; Ethan Rudd; Brian Heflin; Yui-Man Lui; Terry Boult. *Winter Conference on Applications of Computer Vision (WACV 2014)*

2014 *Good Recognition is Non-Metric*
Walter J. Scheirer; **M. Wilber**; Michael Eckmann; Terry Boult. *E. Pattern Recognition 47 (8), 2014*

★ 2013 **Best paper award:** *Animal Recognition in the Mojave Desert: Vision Tools for Field Biologists*
M. Wilber; Walter J. Scheirer; Phil Leitner; et. al.. *Workshop on Applications of Computer Vision (WACV 2013)*

2013 *Issues in Rotational (Non-) Invariance and Image Preprocessing*
Lalit Jain; **M. Wilber**; Terry Boult. *Conference on Computer Vision and Pattern Recognition Biometrics Workshop (CVPRW 2013)*

2012 *PRIVV: Private Remote Iris Authentication with Vaulted Verification*
M. Wilber; Walter J. Scheirer; Terry Boult. *Conference on Computer Vision and Pattern Recognition Biometrics Workshop (CVPR 2012)*

2012 *Secure Remote Matching with Privacy: Scrambled Support Vector Vaulted Verification (S2V3)*
M. Wilber; Terry Boult. *Workshop on Applications of Computer Vision (WACV 2012)*

2011 *Face and Eye Detection on Hard Datasets*
Jon Parris; **M. Wilber**; Brian Heflin; et. al.. *International Joint Conference on Biometrics (IJCB 2011)*